



# Lustre Community Release Update


June 17<sup>th</sup> 2019

Peter Jones, Whamcloud  
OpenSFS Lustre Working Group

# Lustre Community Survey

- Has been running for 8 years
- LWG devises questions
- Useful for tracking trends in Lustre usage
- Full details available at [http://wiki.opensfs.org/Lustre Community Survey](http://wiki.opensfs.org/Lustre_Community_Survey)

Which Lustre versions do you use in production?  
(select all that apply)



Data courtesy of Dustin Leverman (ORNL)


Analysis of commits on git.whamcloud.com v2.11.50 to 2.12.0

# Lustre LTS Releases

- Lustre 2.10.0 GA July 2017
  - Initial LTS branch with regular updates
  - Lustre 2.10.8 GA May 28<sup>th</sup>
    - MOFED 4.6 support
  - Likely the last release on this branch
- Lustre 2.12.0 GA Dec 2018
  - Current LTS branch
  - Lustre 2.12.2 GA May 28<sup>th</sup>
    - MOFED 4.6 support
  - Lustre 2.12.3 targeted for Q3
    - ZFS 0.8 support
    - Tentative RHEL 8.0 client support

# Lustre 2.12


- GA Dec 2018
- OS support
  - RHEL 7.6 servers/clients
  - SLES12 SP3/Ubuntu 18.04 clients
- Interop/upgrades from latest Lustre 2.10.x and 2.11
- Number of useful features
  - DNE Directory Restriping (LU-4684)
  - LNet Network Health (LU-9120)
  - Lazy Size on MDT (LU-9538)
  - T10PI end to end data checksums (LU-10472)
- [http://wiki.lustre.org/Release\\_2.12.0](http://wiki.lustre.org/Release_2.12.0)



# Lustre 2.13

- Targeted for Q3 2019 release
- OS support
  - RHEL 7.6 servers/clients
  - SLES12 SP4/Ubuntu 18.04 clients
  - Tentative RHEL 8.0 clients
- Interop/upgrades from latest Lustre 2.12.x
- Feature Freeze is now in effect
  - Persistent Client Cache (LU-10092) **CONFIRMED**
  - Multi-Rail Routing (LU-11297) **CONFIRMED**
  - Overstriping (LU-9846) **CONFIRMED**
- [http://wiki.lustre.org/Release\\_2.13.0](http://wiki.lustre.org/Release_2.13.0)

# Lustre Community Roadmap



## LEGEND:



### 2.11

- [Data on MDT](#)
- [FLR Delayed Resync](#)
- [Lock Ahead](#)

### 2.12

- [Lazy Size on MDT](#)
- [LNet Health](#)
- [DNE Dir Restriping](#)

### 2.13

- [Persistent Client Cache](#)
- [Multi-Rail Routing](#)
- [Overstriping](#)

### 2.14

- [FLR Erasure Coding](#)
- [Pool Quota](#)
- [DNE Auto Restriping](#)

\* Estimates are not commitments and are provided for informational purposes only

\* Fuller details of features in development are available at <http://wiki.lustre.org/Projects>

# Lustre in Linux Kernel

- Repo for work at <https://github.com/neilbrown/lustre>
  - SUSE and ORNL driving this initiative
  - Plan to submit when code is ready for acceptance
- Major ldiskfs patches merged into upstream ext4/e2fsprogs
  - Now much easier to keep Lustre e2fsprogs current

# Lustre in the Cloud

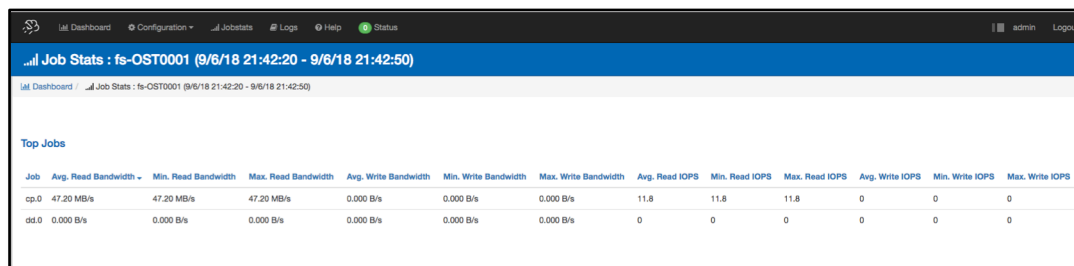
- Intel AWS offering transferred to [DDN-Whamcloud](#)
  - Used in production environments since early 2015
- Amazon announced [FSX for Lustre](#) Software in Nov 2018
  - Presented at [LUG19](#)
- Google have been talking publicly about Lustre
  - [Whamcloud GCP offering](#) launched April 2019
- Multiple offerings on Azure
  - [Cray Azure offering](#) has been available since late 2017
  - [Whamcloud Azure offering](#) launched May 2019



# Integrated Manager for Lustre


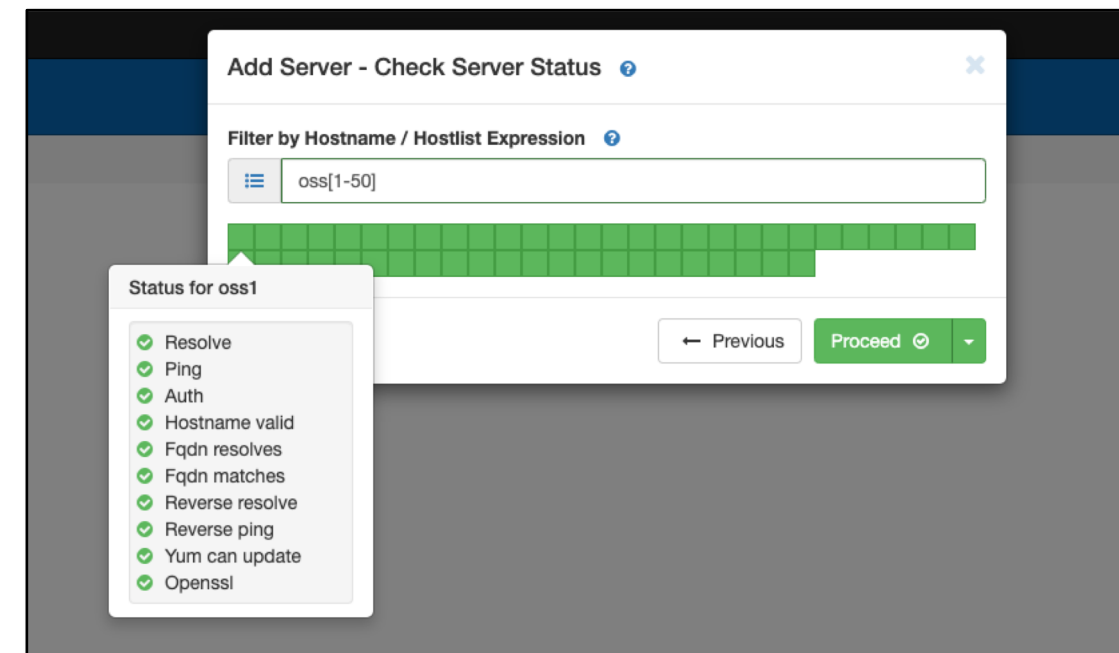
<https://github.com/whamcloud/integrated-manager-for-lustre/releases>

- IML 5.0 now GA
  - Open source management and monitoring tool for Lustre
  - Compatible with Lustre 2.12.x
  - Extending scale of clusters it can manage
  - More flexibility around HA schemes
  - Reduced resource utilization
  - Docker support
  - Simplified install and upgrades



The screenshot shows the IML Dashboard with a navigation bar at the top containing links for Dashboard, Configuration, Jobstats, Logs, Help, and Status. The main content area displays 'Job Stats : fs-OST0001 (9/6/18 21:42:20 - 9/6/18 21:42:50)'. Below this, there is a table titled 'Top Jobs' with columns for Job, Avg. Read Bandwidth, Min. Read Bandwidth, Max. Read Bandwidth, Avg. Write Bandwidth, Min. Write Bandwidth, Max. Write Bandwidth, Avg. Read IOPS, Min. Read IOPS, Max. Read IOPS, Avg. Write IOPS, Min. Write IOPS, and Max. Write IOPS. The table contains two rows of data for jobs 'cp.0' and 'dd.0'.

Job	Avg. Read Bandwidth	Min. Read Bandwidth	Max. Read Bandwidth	Avg. Write Bandwidth	Min. Write Bandwidth	Max. Write Bandwidth	Avg. Read IOPS	Min. Read IOPS	Max. Read IOPS	Avg. Write IOPS	Min. Write IOPS	Max. Write IOPS
cp.0	47.20 MB/s	47.20 MB/s	47.20 MB/s	0.000 B/s	0.000 B/s	0.000 B/s	11.8	11.8	11.8	0	0	0
dd.0	0.000 B/s	0.000 B/s	0.000 B/s	0.000 B/s	0.000 B/s	0.000 B/s	0	0	0	0	0	0



# And Finally...

- **HPC I/O in the Data Center Workshop**
  - Thursday June 20<sup>th</sup> 2pm Frankfurt Marriot Hotel
  - Lustre – the Next 20 Years (Andreas Dilger)  
<https://hps.vi4io.org/events/2019/iodc>



# Thank you

**Open Scalable File Systems, Inc.**

3855 SW 153rd Drive  
Beaverton, OR 97006  
Ph: 503-619-0561  
Fax: 503-644-6708  
[admin@opensfs.org](mailto:admin@opensfs.org)



[www.opensfs.org](http://www.opensfs.org)